arXiv:1406.6176v1 [cs.LG] 24 Jun 2014

Composite Likelihood Estimation for Restricted Boltzmann machines

Muneki Yasuda*, Shun Kataoka, Yuji Waizumi and Kazuyuki Tanaka
*Graduate School of Science and Engineering, Yamagata University, Japan**
Graduate School of Information Sciences, Tohoku University, Japan
*muneki@yz.yamagata-u.ac.jp**

Abstract

Learning the parameters of graphical models using the maximum likelihood estimation is generally hard which requires an approximation. Maximum composite likelihood estimations are statistical approximations of the maximum likelihood estimation which are higher-order generalizations of the maximum pseudo-likelihood estimation. In this paper, we propose a composite likelihood method and investigate its property. Furthermore, we apply our composite likelihood method to restricted Boltzmann machines.

1. Introduction

Learning the parameters of graphical models using maximum likelihood (ML) estimation is generally hard due to the intractability of computing the normalizing constant and its gradients. Maximum pseudo-likelihood (PL) estimation [2] is a statistical approximation of the ML estimation. Unlike the ML estimation, the maximum PL estimation is computationally fast, but however, the estimates obtained by this method are not very accurate.

Composite likelihoods (CLs) [6] are higher-order generalizations of the PL. Asymptotic analysis shows that maximum CL estimation is statistically more efficient than the maximum PL estimation [5]. It has been known that the maximum PL estimation is asymptotically consistent [2]. Like this, the maximum CL estimation is also asymptotically consistent [6]. Furthermore, the maximum CL estimation has an asymptotic variance that is smaller than the maximum PL estimation but larger than ML estimation [5, 3]. Recently, it has been found that the maximum CL estimation corresponds to a block-wise contrastive divergence learning [1].

In the maximum CL estimation, one can freely choose the size of "blocks" which contain several variables, and it is widely believed that by increasing the size of blocks, one can capture more dependency relations in the model and increase the accuracy of the estimates [1]. In the first part of this paper, we introduce a systematic choice of blocks in the maximum CL estimation. In our proposed choice of blocks, it is guaranteed that one can obtain quantitatively closer value to the true likelihood by increasing the size of blocks. In the latter part of this paper, we apply our maximum CL estimation to restricted Boltzmann machines (RBMs) [4] and show results of numerical experiments using synthetic data.

2. Composite Likelihood Estimation

For the n dimensional discrete random variable $\boldsymbol{x} := \{x_i \mid i \in \Omega = \{1, 2, \ldots, n\}\}$, let us consider the probabilistic model expressed as

$$P(\boldsymbol{x} \mid \boldsymbol{\theta}) := Z(\boldsymbol{\theta})^{-1} \exp\big(- E(\boldsymbol{x} \mid \boldsymbol{\theta})\big), \qquad (1)$$

where $E(\boldsymbol{x} \mid \boldsymbol{\theta})$ is the energy function having an arbitrary functional form and $Z(\boldsymbol{\theta}) := \sum_{\boldsymbol{x}} \exp\big(- E(\boldsymbol{x} \mid \boldsymbol{\theta})\big)$ is the normalizing constant. Let us suppose that the data set composed of M data, $\mathcal{D} := \{\boldsymbol{d}^{(\mu)} \mid \mu = 1, 2, \ldots, M\}$ is obtained. Each data is statistically-independent of each other. In the perspective of the ML estimation, we determine the optimal $\boldsymbol{\theta}$ by maximizing the log-likelihood function defined by

$$\mathcal{L}_{\mathrm{ML}}(\boldsymbol{\theta}) := \sum_{\boldsymbol{x}} Q(\boldsymbol{x}) \ln P(\boldsymbol{x} \mid \boldsymbol{\theta}), \qquad (2)$$

where $Q(\boldsymbol{x})$ is the empirical distribution of the data set, i.e. the histogram of data set, expressed by $Q(\boldsymbol{x}) := M^{-1} \sum_{\mu=1}^{M} \delta(\boldsymbol{x}, \boldsymbol{d}^{(\mu)})$, where we define

$$\delta(\boldsymbol{x}, \boldsymbol{d}^{(\mu)}) := \begin{cases} 1 & \boldsymbol{x} = \boldsymbol{d}^{(\mu)} \\ 0 & \boldsymbol{x} \neq \boldsymbol{d}^{(\mu)} \end{cases}.$$

However, maximizing $\mathcal{L}_{\mathrm{ML}}(\boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$ is computationally expensive. This generally requires the computational cost of $O(e^n)$ due to multiple summations.

The maximum CL estimation is a statistical approximation technique of the ML estimation [6]. In the maximum CL estimation, one divides Ω into some different subsets termed *blocks*, $c_1, c_2, \ldots c_r \subseteq \Omega$, with allowing overlaps among blocks. Note that the relation $c_1 \cup c_2 \cup \ldots \cup c_r = \Omega$ must be kept. We denote the family of these blocks, $c_1, c_2, \ldots c_r$, by $\mathcal{F}$. For the family $\mathcal{F}$, the CL is defined by

$$\mathcal{L}_{\mathcal{F}}(\boldsymbol{\theta}) := \Lambda_{\mathcal{F}} \sum_{c \in \mathcal{F}} \sum_{\boldsymbol{x}} Q(\boldsymbol{x}) \ln P(\boldsymbol{x}_c \mid \boldsymbol{x}_{\bar{c}}, \boldsymbol{\theta}), \quad (3)$$

where, for a set $A \subseteq \Omega$, the expression $\boldsymbol{x}_A$ is defined as $\boldsymbol{x}_A := \{x_i \mid i \in A\}$ and $\bar{A} := \Omega \setminus A$. The notation $\Lambda_{\mathcal{F}}$ is defined by $\Lambda_{\mathcal{F}} := |\mathcal{F}|^{-1}$, where the notation $|\cdots|$ denotes the size of the assigned set. From the Bayesian theorem, the conditional probability in the CL is obtained by $P(\boldsymbol{x}_c \mid \boldsymbol{x}_{\bar{c}}, \boldsymbol{\theta}) = P(\boldsymbol{x} \mid \boldsymbol{\theta}) \big(\sum_{\boldsymbol{x}_c} P(\boldsymbol{x} \mid \boldsymbol{\theta}) \big)^{-1}$. In the CL estimation, one maximizes the CL instead of the true log-likelihood. If each block is composed of just one variable, i.e. $c_i = \{i\}$ and $r = n$, the CL is reduced to the PL [2]. Hence, the CL can be regarded as a generalization of the PL. On the other hand, if $r = 1$ and c_1 is composed of all variables, the CL is obviously equivalent to the true log-likelihood $\mathcal{L}_{\mathrm{ML}}(\boldsymbol{\theta})$.

Proposition 1. *The CL generally is an upper bound on the true log-likelihood* $\mathcal{L}_{\mathrm{ML}}(\boldsymbol{\theta})$.

Proof. The relation between original log-likelihood and the CL can be expressed as $\mathcal{L}_{\mathrm{ML}}(\boldsymbol{\theta}) = \mathcal{L}_{\mathcal{F}}(\boldsymbol{\theta}) + \mathcal{R}_{\mathcal{F}}(\boldsymbol{\theta})$, where the remainder term is defined as

$$\mathcal{R}_{\mathcal{F}}(\boldsymbol{\theta}) := \Lambda_{\mathcal{F}} \sum_{c \in \mathcal{F}} \sum_{\boldsymbol{x}} Q(\boldsymbol{x}) \ln \sum_{\boldsymbol{x}_c} P(\boldsymbol{x} \mid \boldsymbol{\theta}). \quad (4)$$

Since $P(\boldsymbol{x} \mid \boldsymbol{\theta})$ is a discrete distribution and $\Lambda_{\mathcal{F}}$ is positive, the remainder term, $\mathcal{R}_{\mathcal{F}}(\boldsymbol{\theta})$, is less than or equal to zero. Therefore, the inequality $\mathcal{L}_{\mathrm{ML}}(\boldsymbol{\theta}) \leq \mathcal{L}_{\mathcal{F}}(\boldsymbol{\theta})$ is generally satisfied for any $\boldsymbol{\theta}$ and for any choice of $\mathcal{F}$. □

2.1. Systematic Choice of Blocks

In this section, we introduce a particular choice of the blocks in which the CL has a good property. For $1 \leq k \leq n$, we define the family $\mathcal{F}_k$ whose elements are all possible blocks composed of k different variables, i.e. $\mathcal{F}_k := \{\{i_1, i_2, \ldots, i_k\} \mid i_1 < i_2 < \cdots < i_k \in \Omega\}$. For example, when $n = 4$, $\mathcal{F}_2 = \{\{1,2\}, \{1,3\}, \{1,4\}, \{2,3\}, \{2,4\}, \{3,4\}\}$ and $\mathcal{F}_3 = \{\{1,2,3\}, \{1,2,4\}, \{1,3,4\}, \{2,3,4\}\}$. For the family $\mathcal{F}_k$, the CL is expressed as

$$\mathcal{L}_{\mathcal{F}_k}(\boldsymbol{\theta}) = \Lambda_{\mathcal{F}_k} \sum_{c \in \mathcal{F}_k} \sum_{\boldsymbol{x}} Q(\boldsymbol{x}) \ln P(\boldsymbol{x}_c \mid \boldsymbol{x}_{\bar{c}}, \boldsymbol{\theta}), \quad (5)$$

where $\Lambda_{\mathcal{F}_k} = |\mathcal{F}_k|^{-1} = k!(n-k)!/n!$. It is noteworthy that $\mathcal{L}_{\mathcal{F}_1}(\boldsymbol{\theta})$ is reduced to the PL and $\mathcal{L}_{\mathcal{F}_n}(\boldsymbol{\theta}) = \mathcal{L}_{\mathrm{ML}}(\boldsymbol{\theta})$.

Proposition 2. *For* $1 \leq k \leq n$*, the CLs for the family* $\mathcal{F}_k$ *is bounded as* $\mathcal{L}_{\mathcal{F}_1}(\boldsymbol{\theta}) \geq \mathcal{L}_{\mathcal{F}_2}(\boldsymbol{\theta}) \geq \cdots \geq \mathcal{L}_{\mathcal{F}_n}(\boldsymbol{\theta}) = \mathcal{L}_{\mathrm{ML}}(\boldsymbol{\theta})$ *for any* $\boldsymbol{\theta}$.

Proof. The relation between $\mathcal{L}_{\mathrm{ML}}(\boldsymbol{\theta})$ and $\mathcal{L}_{\mathcal{F}_k}(\boldsymbol{\theta})$ is $\mathcal{L}_{\mathrm{ML}}(\boldsymbol{\theta}) = \mathcal{L}_{\mathcal{F}_k}(\boldsymbol{\theta}) + \mathcal{R}_{\mathcal{F}_k}(\boldsymbol{\theta})$, where, from equation (4) the remainder term is $\mathcal{R}_{\mathcal{F}_k}(\boldsymbol{\theta}) = \Lambda_{\mathcal{F}_k} \sum_{c \in \mathcal{F}_k} \sum_{\boldsymbol{x}} Q(\boldsymbol{x}) \ln \sum_{\boldsymbol{x}_c} P(\boldsymbol{x} \mid \boldsymbol{\theta})$. Let us consider the difference between the the remainder terms $D_k(\boldsymbol{\theta}) := \mathcal{R}_{\mathcal{F}_{k+1}}(\boldsymbol{\theta}) - \mathcal{R}_{\mathcal{F}_k}(\boldsymbol{\theta})$. After a short manipulation, the difference $D_k(\boldsymbol{\theta})$ yields

$$D_k(\boldsymbol{\theta}) = \frac{\Lambda_{\mathcal{F}_k}}{k-n} \sum_{c \in \mathcal{F}_k} \sum_{i \in \bar{c}} \sum_{\boldsymbol{x}} Q(\boldsymbol{x}) \ln \frac{\sum_{\boldsymbol{x}_c} P(\boldsymbol{x} \mid \boldsymbol{\theta})}{\sum_{\boldsymbol{x}_c, x_i} P(\boldsymbol{x} \mid \boldsymbol{\theta})}.$$

Hence, for $1 \leq k \leq n-1$, the inequality $D_k(\boldsymbol{\theta}) \geq 0$ holds, because $P(\boldsymbol{x} \mid \boldsymbol{\theta})$ is a discrete distribution. Therefore, for $1 \leq k \leq n-1$, the inequality

$$\mathcal{L}_{\mathcal{F}_k}(\boldsymbol{\theta}) \geq \mathcal{L}_{\mathcal{F}_{k+1}}(\boldsymbol{\theta}) \quad (6)$$

is satisfied. From equation (6), we reach to the proposition. □

From propositions 1 and 2, we found that, for $1 \leq k \leq n-1$, the kth-order CL, $\mathcal{L}_{\mathcal{F}_k}(\boldsymbol{\theta})$, is always an upper bound on the true log-likelihood and it monotonically approaches the true log-likelihood with the increase of k. Therefore, it is guaranteed that a larger k gives quantitatively better approximation of the true log-likelihood.

3. Application to Restricted Boltzmann Machines

In this section, we apply the CL estimation to RBMs [4]. RBMs are Boltzmann machines consisting of visible variables, whose states can be observed, and hidden variables, whose states are not specified by observed data. RBMs are defined on (complete) bipartite graphs consisting of two layers. One of them is a layer of visible variables, termed visible layer, and the other one is a layer of hidden variables, termed hidden

layer. There are connections between visible variables and hidden variables, and any interlayer connections are not allowed.

We denote the sets of labels of visible variables and hidden variables by Ω and H, respectively, and we denote the state variable of visible variable $i \in \Omega$ by x_i and the state variable of hidden variable $j \in H$ by h_j. All state variables are binary random variables that take $+1$ or -1. The joint distribution of RBM is represented by

$$P_{\mathrm{RBM}}(\boldsymbol{x}, \boldsymbol{h} \mid \boldsymbol{\theta}) \propto \exp\Big(\sum_{i\in\Omega} \alpha_i x_i + \sum_{j\in H} \beta_j h_j + \sum_{i\in\Omega} \sum_{j\in H} w_{i,j} x_i h_j \Big). \tag{7}$$

The parameters $\boldsymbol{\alpha} = \{\alpha_i \mid i \in \Omega\}$ and $\boldsymbol{\beta} = \{\beta_j \mid j \in H\}$ are biases (or sometimes called thresholds) for visible variables and hidden variables, respectively, and the parameters $\boldsymbol{w} = \{w_{i,j} \mid i \in \Omega,\ j \in H\}$ are weights of connections between the visible variables and the hidden variables. In equation (7), we denote $\boldsymbol{\theta} = \boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \boldsymbol{w}$ for a short description.

Given an empirical distribution $Q(\boldsymbol{x})$ for the visible variables, the log-likelihood of RBM in equation (7) is expressed as

$$\mathcal{L}_{\mathrm{ML}}(\boldsymbol{\theta}) = \sum_{\boldsymbol{x}} Q(\boldsymbol{x}) \ln P_{\mathrm{RBM}}(\boldsymbol{x} \mid \boldsymbol{\theta}), \tag{8}$$

where $P_{\mathrm{RBM}}(\boldsymbol{x} \mid \boldsymbol{\theta})$ is the marginal distribution obtained by $P_{\mathrm{RBM}}(\boldsymbol{x} \mid \boldsymbol{\theta}) = \sum_{\boldsymbol{h}} P_{\mathrm{RBM}}(\boldsymbol{x}, \boldsymbol{h} \mid \boldsymbol{\theta})$. The marginal distribution can be explicitly expressed as $P_{\mathrm{RBM}}(\boldsymbol{x} \mid \boldsymbol{\theta}) \propto \exp\big(-E_{\mathrm{RBM}}(\boldsymbol{x} \mid \boldsymbol{\theta})\big)$, where $E_{\mathrm{RBM}}(\boldsymbol{x} \mid \boldsymbol{\theta}) := -\sum_{i\in\Omega} \alpha_i x_i - \sum_{j\in H} \ln \mathcal{C}_j(\boldsymbol{x} \mid \boldsymbol{\theta})$ and $\mathcal{C}_j(\boldsymbol{x} \mid \boldsymbol{\theta}) := \cosh\big(\beta_j + \sum_{i\in\Omega} w_{i,j} x_i\big)$. The CL estimation can be applied to the RBM. Indeed, the PL estimation for the RBM was introduced [7]. By applying equation (5) to equation (8), we can express the kth-order CL for the RBM as

$$\begin{aligned} &\mathcal{L}_{\mathcal{F}_k}(\boldsymbol{\theta}) \\ &= \Lambda_{\mathcal{F}_k} \sum_{c\in\mathcal{F}_k} \sum_{i\in c} \alpha_i \langle x_i \rangle_{\mathcal{D}} + \sum_{j\in H} \langle \ln \mathcal{C}_j(\boldsymbol{x}, \boldsymbol{\theta}) \rangle_{\mathcal{D}} \\ &\quad - \Lambda_{\mathcal{F}_k} \sum_{c\in\mathcal{F}_k} \Big\langle \ln \sum_{\boldsymbol{x}_c} \exp\big(-E^{(c)}_{\mathrm{RBM}}(\boldsymbol{x} \mid \boldsymbol{\theta})\big) \Big\rangle_{\mathcal{D}}, \end{aligned} \tag{9}$$

where the notation $\langle \cdots \rangle_{\mathcal{D}}$ denotes the average over the empirical distribution and $E^{(c)}_{\mathrm{RBM}}(\boldsymbol{x} \mid \boldsymbol{\theta}) := -\sum_{i\in c} \alpha_i x_i - \sum_{j\in H} \ln \mathcal{C}_j(\boldsymbol{x} \mid \boldsymbol{\theta})$. The gradients, $\Delta^{(k)}_{\theta} := \partial \mathcal{L}_{\mathcal{F}_k}(\boldsymbol{\theta}) / \partial\theta$, with respect to the parameters α_i, β_j and $w_{i,j}$ are

$$\Delta^{(k)}_{\alpha_i} \propto \langle x_i \rangle_{\mathcal{D}} - |\mathcal{F}_k(i)|^{-1} \sum_{c\in\mathcal{F}_k(i)} \langle x_i \rangle_c, \tag{10}$$

$$\Delta^{(k)}_{\beta_j} \propto \langle \mathcal{T}_j(\boldsymbol{x} \mid \boldsymbol{\theta}) \rangle_{\mathcal{D}} - \Lambda_{\mathcal{F}_k} \sum_{c\in\mathcal{F}_k} \langle \mathcal{T}_j(\boldsymbol{x} \mid \boldsymbol{\theta}) \rangle_c \tag{11}$$

and

$$\Delta^{(k)}_{w_{i,j}} \propto \langle x_i \mathcal{T}_j(\boldsymbol{x} \mid \boldsymbol{\theta}) \rangle_{\mathcal{D}} - \Lambda_{\mathcal{F}_k} \sum_{c\in\mathcal{F}_k} \langle x_i \mathcal{T}_j(\boldsymbol{x} \mid \boldsymbol{\theta}) \rangle_c, \tag{12}$$

respectively, where the notation $\mathcal{F}_k(i)$ is the subset of $\mathcal{F}_k$ whose all blocks include i, i.e. $\mathcal{F}_k(i) := \{c \mid i \in c \in \mathcal{F}_k\}$, the notation $\langle \cdots \rangle_c$ is defined as

$$\langle \cdots \rangle_c := \left\langle \frac{\sum_{\boldsymbol{x}_c} (\cdots) \exp\big(-E^{(c)}_{\mathrm{RBM}}(\boldsymbol{x} \mid \boldsymbol{\theta})\big)}{\sum_{\boldsymbol{x}_c} \exp\big(-E^{(c)}_{\mathrm{RBM}}(\boldsymbol{x} \mid \boldsymbol{\theta})\big)} \right\rangle_{\mathcal{D}},$$

for the block c, and $\mathcal{T}_j(\boldsymbol{x} \mid \boldsymbol{\theta}) := \tanh\big(\beta_j + \sum_{i\in\Omega} w_{i,j} x_i\big)$. The computational cost that is required to compute all of them is $O(n^k M |H|)$. Note that, when $k = n$, the gradients (10)–(12) yield the gradients of the true log-likelihood.

3.1. Numerical Experiments

In this section, we show results of numerical experiments using synthetic data. We use an RBM consisting of 5 visible variables and 10 hidden variables as the learning machine, and we generate $M = 70$ data from an RBM consisting of 5 visible variables and 17 hidden variables by using the Markov chain Monte Carlo method. In the generative RBM, we set $\alpha_i = 0.1$, $\beta_j = -0.1$ and $w_{i,j} = 0.2$ for all i and j. We compare the first-, the second- and the third-order CL estimation with the exact ML estimation. We maximize the CLs, i.e. $\mathcal{L}_{\mathcal{F}_1}(\boldsymbol{\theta})$, $\mathcal{L}_{\mathcal{F}_2}(\boldsymbol{\theta})$ and $\mathcal{L}_{\mathcal{F}_3}(\boldsymbol{\theta})$, and the true log-likelihood, $\mathcal{L}_{\mathrm{ML}}(\boldsymbol{\theta})$, by using the gradient ascent method (GAM) with the update rate of 0.1. In the four different GAMs, the same initial values of parameters that are randomly generated are used.

Figure 1 shows that the plot of the CLs shown in equation (9) and the true log-likelihood shown in (8) against the number of iterations of GAMs with the gradients (10)–(12). In this plot, the "ML" is the true log-likelihood obtained by the exact ML estimation and "CL$_k$" are the kth-order CLs, $\mathcal{L}_{\mathcal{F}_k}(\boldsymbol{\theta})$, obtained by the kth-order CL estimations. One can see that the CL approach the true log-likelihood as k increase.

Figure 2 shows the plot of the true log-likelihoods, $\mathcal{L}_{\mathrm{ML}}(\boldsymbol{\theta})$, against the number of iterations of GAMs with

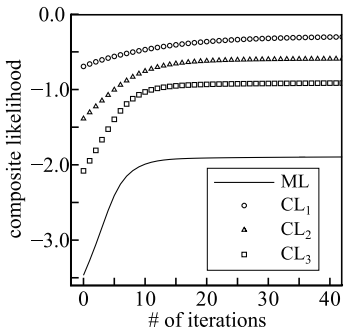

Figure 1. Plot of composite likelihoods against the number of iterations of GAMs. Each point is averaged over 30 trials.

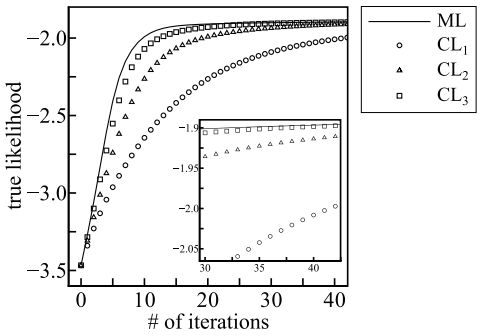

Figure 2. Plot of true log-likelihoods against the number of iterations of GAMs. Each point is averaged over 30 trials.

the gradients (10)–(12). In the plot, the "ML" is the true log-likelihood with the parameters calculated by the exact ML estimation and the "CL$_k$" are the true log-likelihoods with the parameters calculated by kth-order CL estimations. One can see that higher-order CL estimations give better and faster convergence.

After 50000 iterations, the average values (averaged over 30 trials) of the true log-likelihood obtained by the exact ML estimation, the first-, the second- and the third-order CL estimation are -1.741, -1.796, -1.742 and -1.741, respectively. Table 1 shows the mean absolute errors (MADs) of the estimations, $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$ and $\boldsymbol{w}$, between obtained by the exact ML estimation and by the kth-order CL estimation after 50000 iterations. Each MAD is averaged over 30 trials. One can see that higher-order CL estimations give quantitatively better estimations.

Table 1. MADs of estimations between obtained by exact ML estimation and by kth-order CL estimation after 50000 iterations.

	$\boldsymbol{\alpha}$	$\boldsymbol{\beta}$	$\boldsymbol{w}$
$k=1$	0.377	0.431	0.360
$k=2$	0.223	0.223	0.192
$k=3$	0.128	0.114	0.103

4. Conclusion

In this paper, we introduced the systematic choice of blocks for the maximum CL estimation which guarantees that the kth-order CL monotonically approaches the true log-likelihood with the increase of k. This property does not depend on details of graphical models.

Furthermore, we applied our CL method to learning of RBMs and formulate learning algorithm explicitly. In our numerical experiments for synthetic data, we made sure that the higher-order CLs have better performances. As we have seen in section 3, the computational cost increases when higher-order CLs are employed. Nonetheless, it is possible to trade off computation time for increased accuracy by switching to higher-order CLs.

acknowledgment

This work was partly supported by Grants-In-Aid (No. 21700247, No. 22300078 and No. 23500075) for Scientific Research from the Ministry of Education, Culture, Sports, Science and Technology, Japan.

References

[1] A. Asuncion, Q. Liu, A. T. Ihler, and P. Smyth. Learning with blocks: composite likelihood and contrastive divergence. *Proceedings of 13th International Conference on AI and Statistics (AISTAT)*, 9:33–40, 2010.

[2] J. Besag. Statistical analysis of non-lattice data. *Journal of the Royal Statistical Society D (The Statistician)*, 24(3):179–195, 1975.

[3] J. Dillon and G. Lebanon. Stochastic composite likelihood. *Journal of Machine Learning Research*, 11:2597–2633, 2010.

[4] G. E. Hinton. Training products of experts by minimizing contrastive divergence. *Neural Computation*, 8(14):1771–1800, 2002.

[5] P. Liang and M. I. Jordan. An asymptotic analysis of generative, discriminative, and pseudolikelihood estimators. *Proceedings of 25th International Conference on Machine Learning*, pages 584–591, 2008.

[6] B. G. Lindsay. Composite likelihood methods. *Comtemporary Mathematics*, 80(1):221–239, 1988.

[7] B. Marlin, K. Swersky, B. Chen, and N. de Freitas. Inductive principles for restricted boltzmann machine learning. *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 9:509–516, 2010.